

A Method And Apparatus For Transporting Packet Data Over An Optical Network

Inventors: Ping Pan and Ralph Theodore Hofmeister

CROSS-REFERENCE TO RELATED APPLICATION

This application claims priority to U.S. Provisional Application No. 60/440,313, by common inventors, Ping Pan and Ralph Theodore Hofmeister, filed January 15, 2003, and entitled "A METHOD AND APPARATUS FOR TRANSPORTING LAYER-2 TRAFFIC OVER SONET/SDH NETWORKS", which is hereby fully incorporated herein by reference.

BACKGROUND OF THE INVENTION

Field of Invention

The invention generally relates to methods and apparatuses for transporting diverse traffic types such as different types of layer-2 traffic over an optical transport network such as a SONET/SDH network. The invention more particularly relates to utilizing pseudo-wires carried directly on top of the SONET, SDH, or OTN layer to transport diverse data packet traffic types such as various types of layer-2 traffic.

Description of Related Art

Service provider communication networks today consist of multiple types of equipment designed to transmit and route many kinds of traffic. Traditionally, these networks evolved from voice/telephone service so they were designed to carry fixed-sized circuit connections between end users. As data applications have evolved and capacity requirements have grown, several generations of packet switched networking equipment was installed into networks to route the packet data. Examples include ATM, Gigabit Ethernet, and MPLS, as shown in Figure 21.

While new packet switching technologies continue to emerge, service providers must continue to service older technologies as it takes many years for end users to phase

out a particular technology. This has led to the service providers maintaining several independent packet switched networks to carry the different types of service. Provisioning and maintaining these multiple networks is costly it would be advantageous to converge these packet switched networks onto a common network. As shown in Figure 21, Layer-2 and MPLS switches are deployed to aggregate data flows into SONET backbone.

Conventionally circuit switched connections are used to provide transport functions between the various packet switching network equipment. But these circuit switched connections are limited in flexibility: they are available in limited bandwidth sizes: 10Gbps, 2.5Gbps, 622Mbps, 155Mbps, 53Mbps, 1.5Mbps, 64Kbps, and are provisioned and maintained independently of the packet switched traffic. The static nature of these circuit connections imposes inefficiency in utilization of the capacity of the circuit switched network when carrying packet data traffic.

As a result, the interface between the packet data layer (layer 2) of the carrier network and the circuit switch layer (layer 1) leads to network utilization inefficiencies and difficult and expensive provisioning and maintenance tasks for the service providers.

The invention described herein presents a method to couple the Layer-2/MPLS packet data convergence function directly onto circuit switch equipment and integrate the control and management of connections in layer 1 and 2. Integration of these functions will greatly reduce provisioning and maintenance expenses of carrier networks and improve the utilization of the network capacity. The benefit of the invention is evident in Figure 22.

Luca Martini and others have introduced the concept of pseudo-wires in a number of Internet Engineering Task Force (IETF) drafts, which has been widely referred to as "draft-martini". In Martini's design, some pseudo-wires can be initiated from the edge of multi-protocol label switching (MPLS) and/or IP backbone networks. Once established, a customer's layer-2 traffic can be aggregated into the pseudo-wires. To control the pseudo-wires, LDP (label distribution protocol) messages are routed through the backbone network to communicate between network edges. A serious drawback with the draft-martini design is that communication carriers must rely on MPLS/IP backbones with expensive high-performance routers to support the control messaging and label distribution protocol thereby greatly increasing the cost of transporting Layer-2 traffic which is otherwise inexpensive and relatively simple. In

reality, these routers are essentially used to perform relatively trivial switching functionality.

In a parallel development, the Optical Internetworking Forum (OIF) has defined a user-network interface (UNI) specification that allows users to request the creation of Synchronous Optical Network (SONET) connections for data traffic. However, there are a number of issues in the UNI approach:

- Both user and network elements must implement the UNI specification thereby dramatically increasing the cost of implementation and creating compatibility problems with non-UNI networks that interface with the UNI-enabled network.
- The existing OIF UNI is only designed to interface user and network elements over optical interfaces.

George Swallow and others have proposed an overlay model where MPLS routers can use an RSVP (resource reservation protocol extension for traffic engineering) protocol to communicate with a GMPLS-enabled (generalized multi-protocol label switching-enabled) optical backbone. This approach can potentially introduce user traffic aggregation from optical network edges. However, this model requires MPLS and IP to be used across the transport networks. Also, this approach may require the carriers to reveal internal routing and resource information to the external customers, which is not practical in most of the operational networks today.

There have been a number of advancements of SONET/SDH technology in recent years. For example, Virtual Concatenation provides the flexibility that allows edge switches to create SONET/SDH connections with finer granularity bandwidth. Link Capacity Adjustment Scheme (LCAS) uses several control bits in the SONET/SDH frame to increase or decrease a connection's bandwidth. Finally, Generic Framing Procedure (GFP) specifies the framing format for a number of link protocols, such as Ethernet and PPP.

It is admitted that MPLS, LDP, draft-martini, and OIF UNI, Virtual Concatenation, LCAS and GFP are conventional elements with respect to the invention. Although the invention utilizes some of these conventional elements, details of which may be found in available literature, the methods and apparatuses disclosed and claimed herein differ substantially therefrom. In other words, the invention leverages such

conventional technologies in unique ways to achieve a method and apparatus for transporting packet data from customer data nodes over an optical network.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will become more fully understood from the detailed description given herein below and the accompanying drawings which are given by way of illustration only, and thus are not limitative of the present invention, and wherein:

Figure 1 is a network protocol layer model according to the concepts of the invention;

Figure 2 is a simplified network diagram showing a very high level view of the inventive pseudo-wire directly over optical transport network connection techniques according to the invention;

Figure 3 is a network operation model in a high-level block diagram format for explaining network operation according to the invention;

Figure 4 is a structural block diagram illustrating a packet-data-enabled optical connection switch according to the concepts of the invention;

Figure 5 is a functional block diagram illustrating operational details of the inventive packet-data-enabled optical connection switch according to the invention and further illustrating the processing of the data packets on the ingress pathway through the switch;

Figure 6 is a functional block diagram illustrating operational details of the inventive packet-data-enabled optical connection switch according to the invention and further illustrating the processing of the data packets on the egress pathway through the switch;

Figure 7 is a network diagram explaining the operation of control messages according to the concepts of the invention;

Figure 7a is a detailed block diagram illustrating the structure and function of the packet processing engine according to the invention;

Figure 7b is a diagram of the packet filter table structure according to the invention;

Figure 7c is a diagram of the circuit filter table structure according to the invention;

Figure 7d is a diagram of the session table structure according to the invention;

Figure 8 is a high-level block diagram of a packet-data-enabled optical connection switch according to the invention;

Figure 9 is a detailed block diagram of alternative construction and operation of a packet data-enabled optical connection switch and further illustrating an alternative connection of a packet access line module (PALM') according to the invention;

Figure 10 is a high-level block diagram showing an alternative packet-data-enabled optical connection configuration according to the invention and utilizing the alternative packet access line module of Figure 9;

Figure 11 is a high-level block diagram showing one alternative data flow within the packet-data enabled optical connection switch configuration of Figure 10;

Figure 12 is a high-level block diagram showing a second alternative data flow within the packet-data-enabled optical connection switch configuration of Figure 10;

Figure 13 is a high-level block diagram showing a third alternative data flow within the packet-data-enabled optical connection switch configuration of Figure 10;

Figure 14 is a high level flowchart illustrating the general operation of the invention from both the transmit and receive perspectives.

Figure 15a is a flow chart illustrating the inventive processing of a data packet received from a data port;

Figure 15b is a flow chart illustrating the inventive processing of a packet fetched from an optical connection including the processing of both data packets and control messages;

Figure 16 is a flow chart illustrating the inventive method of injecting a control message into an optical interface;

Figure 17 is a sequence diagram showing the inventive method of setting up data flow over an optical connection;

Figure 18 is a sequence diagram showing the inventive method of removing a data flow over an optical connection;

Figure 19 is a sequence diagram showing the inventive method of handling the situation in which the optical connection has failed or become deactivated;

Figure 20 is an example of the inventive apparatus and methods in operation;

Figure 21 is a model of a conventional network used by communication providers; and

Figure 22 is a model of the network according to the principles of the invention.

DETAILED DESCRIPTION OF INVENTION

The following detailed description of the invention refers to the accompanying drawings. The same reference numbers in different drawings identify the same or similar elements. Also, the following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims and equivalents thereof.

The expression "optically communicates" as used herein refers to any connection, coupling, link or the like by which optical signals carried by one optical system element are imparted to the "communicating" element. Such "optically communicating" devices are not necessarily directly connected to one another and may be separated by intermediate optical components or devices. Likewise, the expressions "connection" and "operative connection" as used herein are relative terms and do not require a direct physical connection.

Definitions:

The invention described below utilizes various terms that may or may not be fully consistent with the conventional understanding of similar or identical terms. To clarify the meaning of these various terms the following definitions are used by this invention description:

a)MAC: media access control: The interface to the physical media. Assembles and disassembles frames and controls physical interface communications. The physical interface and frame format is L2-specific so that different client interfaces will contain specific MAC devices and/or multi-purpose MAC devices.

b) PALM: Packet Access Line Module. unit that originates and terminates packet data traffic from/to other equipment via physical interfaces. The PALM differs from the TDM (Time Division Multiplexed) Line Module in that it terminates packet data physical interfaces and frames and processes the packet traffic. The PALM described in more detail below generally contains the PPE (packet processing engine) and PPE controller and performs the translation and aggregation of packet data to/from optical connections. The simplified PALM' in the server architecture does not

originate/terminate the optical connections. Instead, it translates packet data to/from internal connections between the PALM' and the server cards.

c) PPE: Packet Processing Engine. Performs per-packet forwarding decisions, appends/removes encapsulation labels, processes and delivers control messages, collects performance statistics, polices incoming traffic and shapes outgoing traffic.

d) optical circuit switch: A network element that switches and manages optical connections.

e) line module: a field-replaceable unit of the switch that contains the physical ports for traffic termination and origination.

f) data flow: a sequence of data packets that are associated with one another. All packets in a flow originate at the same node and terminate at the same node but not all packets with the same origin and termination are necessarily in the same flow as one another.

i) customer data flow: includes all types of L2 and MPLS packets. Flows from/to the client edge are differentiated by one another by various means, depending on the physical interface and frame format of the data link layer.

ii) provider data flow: also feeds into the line modules being used as well as the various node definitions below. The invention does not depend on the topology or protection scheme of the optical network. The invention simply requires a point-to-point connection between two provider edge nodes.

g) provider edge node: the nodes at which client data packets from a flow are translated from/to an optical connection. Packets in a flow will traverse two and only two provider edge nodes: the ingress and the egress.

h) customer edge node: the node originating (terminating) the data link layer session terminating (originating) on the provider edge node client port.

i) intermediate provider nodes: nodes that the optical connection traffic passes through between the ingress and egress provider edge nodes. The intermediate nodes do not have to be aware of the data flows contained within the optical connections. Their primary

function is to switch/manage the optical connection as they would a traditional or non-data flow optical connection.

j) encapsulation label: A unique identifier contained in every data packet traversing the optical connection, used to differentiate pseudo-wires. The encapsulation label is normally appended by the ingress provider edge node and removed by the egress provider edge node. However, it is possible and may be desirable in some cases for the encapsulation label to be appended and/or removed by a customer node, or over-written by an intermediate provider node.

k) pseudo-wire: a logical point-to-point connection between two provider edge nodes that is used to forward data packets from one and only one flow. One or more pseudo wires may be contained in an optical connection. A pseudo wire differs from a flow in that: 1) it originates and terminates on provider edge nodes while a flow does so on customer nodes; 2) the arrival sequence of packets will be maintained over the pseudo wire while a flow may not guarantee the sequence of packets.

l) control message label: a unique label such as the IP4 Explicit NULL Label that distinguishes control messages from data packets. In general, a unique encapsulation label to differentiate packets in an optical connection that are used by the provider nodes to pass management and control information between themselves.

m) control message: a message or signal that is used to control the provider network, customer edge nodes, components thereof, or the data being transported across the provider network or to the customer edge nodes. The invention does not generate the control messages or effect control based on them. Instead, the invention is concerned with transporting such conventional control messages. In general, the invention can practically tunnel any appropriate control message. Some illustrative but non-limiting examples are as follows:

1. control messages relating to MPLS/IP control protocols: such control messages are used to discover and establish pseudo-wires as well as MPLS labeled-switch-path. All of these MPLS/IP control messages may be aggregated into an optical connection with label Explicit-Null or other control

message encapsulation label according to the invention as discussed in detail below. Some of the more important categories of control messages may be taken from the following protocols: LDP (label distribution protocol), RSVP (resource reservation protocol), and OSPF (open shortest path first).

2. IP Data control messages: To ensure the connectivity between two edge nodes, the user can aggregate probing packets from an edge node, and check if they can be received at the other edge node. Such probing packets are defined in ICMP (internet control message protocol) and LSP-ping (a special sequence of packets designed to detect the connectivity of MPLS LSPs as known in the art..

3. Layer-2 messages: To interconnect two layer-2 data interfaces through an optical connection, it is possible to tunnel conventional Layer-2 control messages such as ARP (address resolution protocol) PAUSE (a signaling protocol in Ethernet for flow control), heartbeat messages between two nodes through an identifiable control message encapsulation label according to the teachings of the invention.

4. Control messages relating to upper application data: when supporting IP encapsulated packets, such as real-time traffic using RTP (real time protocol) which are used to convert real-time streams into IP packets. The invention can pick out or capture the in-band control packets within RTP packets such as RTCP (Real Time Control Protocol) packets and deliver them to the other edge of the optical connection. This will allow the edge nodes to monitor real-time flows, and enforce associated QoS for the flows.

General Description

In general terms, the invention initiates and maintains pseudo-wires directly over existing SONET networks using the already-deployed SONET switching gear. In the invention, unlike a UNI-based network, the switching intelligence only needs to be implemented in the SONET switches (network elements) and the users are not required to implement additional functionality. Furthermore, the invention works over a wide

variety of customer interfaces including Ethernet, ATM, and Frame Relay optical and/or copper interfaces.

By examining some of the existing communication backbone topologies and traffic patterns, the inventors noticed that much of the data traffic comes from traditional switching networks: Ethernet, Frame Relay and ATM. Typically, voice traffic can be transported via Frame Relay circuits, and ADSL is based on ATM. With the recent rapid advancement in Gigabit Ethernet technology, Ethernet interfaces have been gradually deployed at places where both IP and non-IP traffic aggregation takes place.

Hence, the invention represents a very practical application that enables carriers to “tunnel” user traffic through well-provisioned SONET transport backbones from the edge of their networks. Further, the idea of developing yet another layer of tunnels on top of SONET cross-connections, such as building MPLS LSPs (label switched paths) as is being proposed by router vendors, is not economically practical or technically beneficial.

The invention creates “pseudo-wires” over, for example, SONET cross-connections directly, and switches layer-2 MAC frames from network edges, reducing cost and complexity of the network switching elements. The invention may utilize many of the conventional mechanisms for setting up pseudo-wires but in unique ways as explained herein. Details of the conventional pseudo-wire mechanisms are well known and need not be discussed here in detail. Instead, this disclosure focuses on the adaptation of pseudo-wire techniques such that a pseudo wire may be carried directly over a provisioned SONET network. Alternatively, the pseudo wire may be carried directly over a provisioned Synchronous Digital Hierarchy (SDH) or Optical Transport Network (OTN) network.

The inventive protocol-layering model is shown in Figure 1. It is important to realize that this protocol-layering model is different from the current framework, where pseudo-wires are created on top of either MPLS or IP GRE (generic routing encapsulation) tunnels which are, in turn, carried on top of the SONET transport layer.

One constraint in the conventional framework is that to create and manage MPLS or GRE tunnels (generic routing encapsulation), IP routing, IGP (interior gateway protocols) and BGP (border gateway protocol) and signaling RSVP-TE (resource reservation protocol extension for traffic engineering) and LDP (label

distribution protocol) have to be used throughout the network. Therefore, to transfer layer-2 traffic according to conventional schemes such as those proposed by Luca Martini, the carriers have to rely on an IP overlay network between the layer-2 switching networks and the transport networks. Due to backbone traffic volume, high-end expensive backbone routers are required to construct such overlay networks. This design could result in adding tremendous cost to carriers, while their existing SONET transport links and equipment may be under-utilized. Also, maintaining an additional overlay IP network increases the network management and operation cost to the carriers.

Thus, to achieve the objectives of transporting layer-2 traffic, the inventors create pseudo-wires over, for example, SONET cross-connections directly, and support draft-martini (or equivalent) on SONET switches at network edges to setup and manage pseudo-wires. No router over-layer network is required in the inventive design.

Returning to Fig. 1, the protocol-layering model includes the conventional SONET transport layer that creates and maintains SONET cross connections in the conventional fashion. The pseudo-wiring is carried directly on top of the SONET transport layer according to the inventive protocol-layering model. Such pseudo-wires may be used to carry Layer-2 traffic such as Ethernet MAC, ATM, Frame Relay, etc. as well as MPLS data packets. In general, any packetized traffic may be carried by the pseudo-wires. The next layer is the actual payload which may be any data including voice, data packets, etc as is well known in the art.

It is important to realize here that, in the conventional model proposed in IETF and Luca Martini, the pseudo-wiring layer situates above IP layer. Below the IP layer is MPLS, Layer-2 and transport layers, respectively. One of the main reasons for such a model is to use IP layer for control message delivery. Since only routers have the ability to deliver control messages through the Internet backbone, pseudo-wiring therefore becomes a router-only application. In contrast, the invention utilizes the conventional SONET transport layer to deliver control messages between edge nodes. As a result, pseudo-wiring can be accomplished on devices other than routers at a much cheaper cost.

Overview Of Operation

Before proceeding to the apparatus details, a general overview of the inventive operation is provided. Setting up pseudo wires (PW) may follow a procedure as defined in [PWE3-CTRL (L. Martini, et al, "Transport of Layer 2 Frame Over MPLS", draft-ietf-pwe3-control-protocol-05.txt)], but this procedure is modified by the invention to operate in the context of PW directly on top of the SONET, SDH, OTN or equivalent layer. The operation reference model for a SONET system is shown in Figure 2 but it is to be understood that substantially the same reference model applies for SDH and OTN.

As shown in Fig. 2, the inventive network includes customer data nodes such as customer data nodes 1 and 2. A customer data node may be a conventional switch or router. The provider edge node generally includes conventional SONET cross-connect functionality but implemented by a data-enabled SONET switch according to the invention such as the one illustrated in Fig. 4 and further explained below.

From the customer network edge (customer data nodes as illustrated in Figure 2 represent the customer network edge), data flow such as layer-2 frames enter the provider's backbone. More specifically, a data packet such as a layer-2 frame may be sent from customer data node 1 to provider edge node 1. The provider edge nodes 1, 2 set up a SONET cross-connection in the usual and conventional fashion across the provider network.

The invention then sets up a pseudo wire directly within the SONET cross-connect as further illustrated in Fig. 2. The pseudo-wire and SONET cross-connect are terminated at the other end of the provider network, in this case at provider edge node 2. The provider edge node 2 then transmits the data flow (e.g. layer-2 frames) to the customer data node 2.

It is to be understood that the provider network typically includes far more than 2 edge nodes and that intermediate nodes are also typically included but for ease of illustration such additional nodes are omitted from Fig. 2.

Each of the layer-2 frames within the layer 2 flow has a "flow-id" in their header. The flow-id may be implemented with conventional VLAN-ID's for Ethernet MAC frames, DLCI for Frame Relay frames, and VPI/VCI for ATM cells. It is also a possibility that the customer edge equipment may inject MPLS frames into the backbone. The use of this flow-id for the setting up and maintenance of pseudo wires according to the invention is further explained below.

Figure 3 is a network operation model according to the invention and is useful for illustrating the general concepts of the invention. The customer data nodes (A, E) and provider edge nodes (B, D) may be implemented as discussed above. The backbone network is a conventional optical network such as a SONET, SDH or OTN-based network that is typically part of a provider network.

In reference to Figure 3, data packets travel from A to E through B, C and D. Each packet is encapsulated with either Layer-2 and/or MPLS label. Each Layer-2 and MPLS label uniquely identifies one data flow between two nodes. In this description, when such a data flow is placed onto the provider network according to the inventive teachings it is referred to as "pseudo-wire". Further, it is assumed that the data flows and pseudo-wires are bidirectional, although the mechanism defined here does not exclude the operation for uni-directional traffic. It is further assumed that the backbone network, C, is a conventional carrier's transport network utilizing conventional mechanisms such as SONET-switching to deliver data. In other words, no modifications are necessary for the backbone network C elements to carry the inventive pseudo-wires.

Provider edge nodes B and D are the devices to which this invention will apply and represent the network elements that would be modified (or replaced) according to the invention. Provider edge nodes B and D are capable of performing both data switching and circuit switching. "Data switching" means that the packets are forwarded based on Layer-2 and MPLS headers. "Circuit switching" means all data sent to the circuit will be routed through the network along the same path from the time the circuit is established until it is terminated.

Upon the completion of inspecting an incoming data packet, provider edge nodes B and D will encapsulate the data packet with a label that can uniquely identify the user flow to which the packet belongs, and send the packet to a pre-established circuit over backbone network C. At egress, provider edge nodes B and D will recover the packet from the circuit, remove the label and transmit the packet out to the proper destination. There exist one or multiple circuits between provider edge nodes B and D. Each circuit can aggregate one or multiple pseudo-wires.

From the control plane perspective, it takes two steps to initiate a pseudo-wire over a circuit between provider edge nodes B and D. The first step requires the network operator, F, to download the mapping between the pseudo-wires and the circuits to the provider edge nodes B and D. The creation of the mappings may be the

result of a prior business agreement, or bilateral agreement between carriers, and is beyond the scope of this invention.

Once the mapping information has been received and processed on provider edge nodes B and D, B and D will start to negotiate with each other to agree upon the encapsulation labels that pseudo-wires should use for packet encapsulation. By default, provider edge nodes B and D will allocate two encapsulation labels for each pseudo-wire, one for receiving and another for transmitting. Upon the completion of the label negotiation, provider edge nodes B and D will update the encapsulation label information to the data plane, and thus a pseudo-wire has been created.

At any given time, provider edge nodes B and D may inform operation status to network operator F. Likewise, network operator F may query provider edge nodes B and D for control and accounting information. However, it is beyond the scope of this invention to further specify the relationship between network operator F and customer (client) data nodes, A and E.

The apparatus elements within the provider edge nodes that is responsible for the functionality described above is shown in block diagram form in Figure 4. As shown therein, the inventive modifications are within an optical circuit switch such as a SONET, SDH or OTN optical circuit switch and transform the conventional optical circuit switch into what is termed herein a "packet-data-enabled optical connection switch" which is represented as element 5 in the drawings.

As shown in Figure 4, the packet-data-enabled optical connection switch 5 includes a packet access line module (PALM) 10 that receives packet data from a port. This is diagrammatically indicated by a data flow arrow but the physical port will also include an appropriate physical interface (not shown) the conventional construction of which will vary depending upon the type of packet data being received and physical interface (optical, copper, line rate) as is known in the art. The PALM 10 is operatively connected to a TDM switch fabric 30 which may be constructed with a known cross connecting TDM switch fabric such as those used in conventional SONET switches one example of which is used by the CoreDirector® switch made by CIENA Corporation.

Figure 4 is a simplified drawing for the purposes of explaining the processing of a single data flow and therefore shows only a single PALM 10 having only one port receiving a data flow. Likewise, the simplified drawing of Fig. 4 only shows one output from the TDM switch fabric to a single TDM line module 40. It is to be

understood that the actual implementation would have a plurality of ports for the PALM 10. Furthermore, the actual implementation would have a plurality of ports that feed into the TDM switch fabric 30 and that the TDM switch fabric 30 output will feed into a plurality of TDM line modules 40 and output ports. In addition, it is to be understood that the implementation would have a mechanism to aggregate packet data from a plurality of PALMs prior to transmitting into the optical connections. Such a mechanism for aggregating packet data is known in the packet data switch art and could be included in the inventive packet-data-enabled optical connection switch 5, 5'.

In general, the packet fabric 34 provides connectivity between PPEs and the PPEs perform the aggregation. Even without aggregation over multiple PALMs there could still be other types of aggregation performed by the invention because a single PALM 10 may have multiple physical ports and flows from different ports may be mapped to pseudo-wires that reside in a common optical connection.

Examples of a full packet-data enabled optical switch are explained below in reference to Figures 8 and 10.

The TDM line module(s) 40 are conventional elements in and of themselves and provide the functions of framing (via conventional framer 45 included therein) and, electrical-to-optical conversion, and optical signal transmitting such that the data may be carried as an optical signal over the provider network. The framer 45 is a very conventional element and may utilize conventional optical transport framing schemes such as SONET, SDH, OTN, or a future developed optical signal transport framing scheme. It is greatly preferred that standardized optical transport framing schemes be used so as to take advantage of and otherwise leverage the existing optical networks utilizing such standardized framing schemes. In the U.S., this would mean SONET while in Europe it would be SDH since those are the respectively prevailing standards at this time.

The PALM 10 includes a media access controller (MAC) 12 which is a conventional element receiving packet data and terminating the customer data flow. The MAC also extracts the packet data such as an L2 packet from the customer data flow. The MAC 12 is connected to a packet processing engine (PPE) 15 that is a unique element constructed according to the principles of the invention as further discussed below in relation to Figures 7a-d.

The packet processing engine 15 has access to a mapping database 19-1 that contains mapping tables (packet filter table 60 subset and circuit filter table 80 subset which are explained below in relation to Figures 7a-c). The PPE 15 also has access to a control message database 18-1 which includes a session table 25 subset. Generally speaking, the PPE 15 classifies the incoming packet or otherwise determines what type of packet is incoming, polices the data flow, collects performance statistics, appends an appropriate encapsulation label, aggregates traffic and shapes the outgoing traffic for logical circuits. Aggregation of traffic is possible since a single optical connection (e.g. a subnetwork connection which may be at a rate of OC-12, OC-48, etc) may hold more than one pseudo wire containing a packet. Further details of the PPE 15 operation are provided below in relation to Fig. 7a.

The PPE is operatively connected to the mapping engine 17 which is itself a conventional element that encapsulates the packet + label. One example of such encapsulation that may be used by the invention is the conventional GFP (Generic Framing Procedure as defined by ITU-T G.7041/Y.1303). Other examples include LAPS (Link Access Procedure – SDH, ITU standard X.86), PoS (Packet over SONET IETF RFC2615) and other HDLC-framing methods (such as the ones used on Cisco routers).

The mapping engine 17 also originates and terminates optical connections as is known in the art (e.g. optical connections using SONET, SDH or OTN). The mapping engine, in one implementation originates/terminates the optical connection. The TDM fabric 30 and TDM LM/framer 45 allow muxing/demuxing of the optical connection so that it may go out one or more physical ports and share the physical port with other TDM traffic and/or other PW-carrying optical connections. These optical connections output from the mapping engine are then sent to the TDM switch fabric 30 that switches the connections (or circuit elements if virtual concatenation is used). The switch fabric 30 is connected to a TDM line module 40 which includes a framer 45 that implements a conventional SONET, SDH, or OTN optical transport framing and originates/terminates the optical transport signal to/from the provider network.

As mentioned above and as shown in Fig. 4, the main data flow pathway through the packet-data-enabled optical circuit switch 5 is a bidirectional pathway. Although the above description mainly focuses on the left-to-right (ingress) flow taking the customer data flow and processing it to output an optical signal on the provider network, the

reverse (egress) flow is also part of the invention. This is further discussed below in relation to, for example, Figs. 5 (ingress flow) and 6 (egress flow).

As further shown in Figure 4, a switch controller 20 has control over the MAC 12, PPE 15, mapping engine 17, TDM switch fabric 30 and TDM line module 40. The switch controller 20 may be constructed with, for example, a general-purpose microprocessor and associated memory, ASIC(s), FPGA(s) or other well-known techniques for building such control modules. The control functions performed by the switch controller 20 are programmed into the microprocessor, ASIC, FPGA, etc. Conventional aspects of switch controller 20 functionality such as certain conventional aspects of control over the TDM switch fabric 30, line module 40, MAC 12 and mapping engine 17 are not described in detail herein. As appropriate, this disclosure focuses on the novel aspects of control exercised by the switch controller 20 and are explained in detail below particularly in relation to Figs. 17-19. Generally speaking, the PW label negotiation is performed by the switch controller 20 as the PPE 15 typically cannot provide system-wide label allocation and network view etc. Once the labels have been negotiated, the switch controller 20 will download the negotiated labels to the PPEs 15 for data switching. The switch controller 20 has access to a control message database (DB) 18 which includes a session table 25. The database 18 holding session table 25 may be stored in a separate memory module as shown in Fig. 4 or it may be stored in a common memory module along with the mapping tables of database 19. More specifically, the switch controller 20 maintains a master copy of all information including a master copy of the control message database 18 storing the session table 25 and a master copy of the mapping database 19 including the packet filter table 60 and circuit filter table 80. The switch controller 20 distributes the information from all of these tables to the PPE 15 on each individual PALM 10.

In a full packet-data-enabled optical connection switch 5 such as the one shown in Fig. 8, the switch controller 20 controls a plurality of PALMs 10-1 through 10-n each of which includes a PPE 15. Continuing this notation, the individual PPEs 15-1 through 15-n each have a corresponding subset of the control message database 18 and the mapping database 19. Thus, the individual PPEs 15-n each store a control message DB subset 18-n (storing a session table 25 subset) and a mapping DB subset 19-n (storing a packet filter table 60 subset and a circuit filter table 80 subset).

Figure 5 further illustrates the inventive ingress processing of packet data arriving as a client signal. In detail, Figure 5 shows the main elements of the packet-

data-enabled optical connection switch 5 including MAC 12, PPE 15, mapping engine 17, TDM switch fabric 30 and framer 45. A customer data flow arriving at the MAC 12 may be in a wide variety of formats including but not limited to GE (gigabit Ethernet), LAPS (link access procedure - SDH), EoS (Ethernet over SONET), ATM (asynchronous transfer mode), FR (frame relay), RPR (Resilient Packet Ring IEEE 802.17), POS (Packet over SONET) or any layer 2 packet with or without an MPLS label.

All of these data types are represented in Fig. 5 as a layer-2 packet (L2 pkt) after the associated transport frame structure has been removed. As shown therein, the MAC 12 extracts the L2 packet. The PPE 15 appends an appropriate encapsulation label (further discussed below) which is shown as "L2 pkt/Label" in Fig. 5. The mapping engine encapsulates the L2 packet with the encapsulation label in a GFP frame or equivalent and optical connection frame structure. The mapping engine further encapsulates the packet as necessary in a compatible format for the TDM switch fabric 30. The packet traverses the TDM switch fabric in the optical connection to one or more framers 45 where the optical connection may be groomed with other optical connections and prepared for transmission in a conventional optical frame such as a SONET frame, SDH frame or OTN frame.

Figure 6 further illustrates the reverse or egress path through the packet-data-enabled optical connection switch 5 from the perspective of the data flow through the packet-data-enabled optical connection switch 5. Specifically, data packets transmitted through the optical transport network via pseudo-wires carried on optical connections are received by the framer 45 where the underlying SONET, SDH, or OTN frame structure is terminated and the payload envelope is converted as necessary into a compatible format for the TDM switch fabric 30. The data packets traverse the TDM switch fabric to the mapping engine 17, which converts as necessary from the TDM switch fabric format, terminates individual optical connections, and extracts packets and removes the GFP or equivalent frame overhead. The underlying packet that still includes the encapsulation label is passed to the PPE 15. The PPE determines the appropriate physical port to send the packet out on and optionally overwrites the L2 label based on the encapsulation label value and the optical connection it was received on. The PPE removes the encapsulation label prior to passing the L2 packet to the MAC. The MAC encapsulates the L2 packet in the appropriate L1 frame/format and sends it to the physical port for transmission to the customer edge node.

Edge-to-Edge Message Tunneling

Figure 7 illustrates the structure of a network that can aggregate multiple data flows over a single optical connection. There exists an optical connection between two Packet-Data-Enabled Optical Connection Switches, C and H that are built according to the invention (e.g. the packet-data-enabled optical connection switch 5, 5' as described herein). The remainder of the nodes A, B, D-G, I and J are conventional equipment. The optical connection can be in the form of, for example, a SONET, SDH or OTN transport circuit.

The optical connection can aggregate multiple data flows from Customer Nodes A, B, I, and J. Each flow is associated with a unique encapsulation label at either receive or transmit direction. The packets that belong to a particular flow will be encapsulated with a label at C and H. The value of the label is the result of control-plane negotiation between C and H as further explained below.

One critical issue in this architecture is the delivery of the control messages. Obviously, to support large number of data flows, each Data-Enabled Optical Switch may require processing a large volume of control traffic. There are a number of methods to accomplish this including:

1. Route control messages through the network. This is the method used in the Internet, where each control packet is delivered hop-by-hop until it reaches to the final destination. Note: in the similar method of aggregating data flows over MPLS network [draft-martini], the control packets are "routed" through the router network. This approach is not practical in optical networks, since this would require every optical node to establish a special connection to a neighboring optical node for the purpose of delivering control messages only.
2. Send control messages through SONET DCC channel: the DCC channel is a set of control overhead fields in SONET frames. It has been used to exchange control messages between optical nodes within optical networks. DCC channels, however, have very limited bandwidth. The option of inserting data-control messages to DCC channels may cause traffic congestion which would result in optical network internal information loss.

3. Out-of-band signaling: Like SS7 networks operated in PSTN networks, one option is to build an out-of-band control network for control message delivery. However, this can be very costly in terms of network manageability.

After evaluating all the existing options, the inventors created an in-band method for control message delivery. The idea is to treat control messages as regular data packets, and inject them into the optical connection that they are supposed to provision for the data flows. In other words, in the invention, all control packets are to be “tunneled” through SONET (or SDH or OTN) cross-connections as regular payload from the edge. Each data flow is associated with a label, and the invention encapsulates each control message with an identifiable encapsulation label that can be recognized by the edge nodes.

In Figure 7, there exists an optical connection going through nodes C, D, E, G and H. The provider edge nodes D and H include a data enabled optical switch 5 according to the invention such that C and H will use the connection to exchange control messages. Each control message is encapsulated with a label that both C and H can recognize. Subsequently, C and H will capture and send the control messages to the control plane for processing. One example of an identifiable label is the Explicit NULL label defined in Rosen et. al, “MPLS Label Stack Encoding”, RFC3032, Network Working Group, Request for comments 3032 submitted to Internet Society, January 2001 which may be found at <http://www.ietf.org/rfc/rfc3032.txt>). The identifiable label is also called a control message encapsulation label herein and is not limited to the NULL label mentioned above. Indeed, any label could be used as the control message encapsulation label. For example, the provider edge nodes may negotiate any label to serve as the control message encapsulation label and such a label will thereafter identify the data packet as a control message.

There are a number of advantages in the inventive approach described herein including:

1. Control message processing only involves the edge nodes. Network intermediate nodes are not disturbed, need no modification and merely pass along optical signals in the normal fashion. In Figure 7, other than the provider edge nodes C and H, the rest of the optical nodes (D, E, F, and G) are not aware of the existence of control messages.

2. Since control messages are encapsulated with labels, this simplifies the processing overhead at the provider edge nodes. The control messages are processed as regular data packets. Instead of sending out to a data interface, they are forwarded to the control module. The detailed mechanism for accomplishing this is elaborated upon below.

3. Since control messages traverse the same optical connections that data flows will traverse, it is easier and faster for the edge nodes to react to network failures. In comparison, in MPLS networks, when there is a failure on the data plane, it will take seconds before the control plane will be aware of the problem - likely to be notified from the routing protocol updates. In the inventive approach, the control-plane and the data-plane share the same fate. As a result, the control-plane can respond to failures faster. This is a huge advantage particularly because protection mechanisms can be triggered much faster thereby preventing data loss. At modern line rates currently approaching 40 gigabits/seconds per wavelength activating protection mechanisms in a shorter time will prevent the loss of tremendous amounts of data.

Generally speaking, the invention operates as follows. When a data flow such as a layer-2 frame is received from a user's network, the PPE 15 encapsulates (or pushes) a pre-negotiated encapsulation label onto the packet. On the other hand, when a control packet (such as LDP Hello message) needs to be delivered through the network, the invention pushes an identifiable label such as the "IP4 Explicit NULL Label" on to the control message. The PPE 15 will direct all frames into the pre-established SONET connections (pseudo-wires). Further detailed operation is provided below in relation to Figures 14 and 16.

On the other end of the SONET connection, the PPE 15 will de-encapsulate (or pop) all received frames. For data packets, the PPE 15 forwards them to the user network. If the received label is the identifiable control message label (e.g. "IP4 Explicit NULL Label"), the PPE 15 forwards the message to the switch's central processor 20 for further processing. Further details are provided below in relation to Figures 14 and 15b.

Figure 14 is a high level flowchart illustrating the general operation of the invention from both the transmit and receive perspectives. All of the operations outlined in Fig. 14 are performed by the PPE 15. As shown therein, the invention first establishes (300) an optical connection between two provider edge nodes which is a conventional process in and of itself that may use conventional SONET, SDH or OTN techniques to do so. The data packets may be aggregated into this optical connection. Next, the PPE 15 tunnels (305) packet data within the established optical connection. Pseudo-wires may then be established (310) by tunneling command messages within the same established optical connection as used for the packet data. Like the data packets, the control messages may also be aggregated within the same optical connection at least to the extent the control messages share the same optical connection pathway through the provider network. When transporting control messages, the PPE utilizes (320) a distinguishable encapsulation label for the command message. Such a distinguishable encapsulation label is also referred to herein as a control message encapsulation message.

On the receive end, as further shown in Fig. 14, the PPE 15 parses the encapsulation label from the received data. The PPE 15 may then decide (330) whether the parsed encapsulation label matches the command encapsulation label type. If yes, then the received message is processed (340) as a command message a process which may include sending the command message to the switch controller 20. If the parsed label does not match the command message encapsulation label type, then the received message is processed (335) as a data packet a process which may include using the parsed label to lookup the outgoing data interface from the circuit table 80 that applies to the particular data packet just received.

Figure 7a is a detailed block diagram of the packet processing engine (PPE) 15 that is a key part of the invention and which may, for example, be part of the packet access line module 10 as shown in Fig. 4.

The packet processing engine 15 is the device responsible for processing incoming data packets, mapping packets into optical connections, processing packets received from optical connections, and injecting control messages into optical connections. Unlike traditional switching devices that perform either packet or circuit switching, in the invention design, each PPE 15 operates for both packet and circuit switching simultaneously.

The processing of data packets includes operations such as packet header lookup, extra header encapsulation, and packet switching into optical connections. The processing of packets from optical connections includes operations such as SONET Path Over Head (POH), packet header manipulation and label switching. One SONET POH handling is the ability to work with Virtual Concatenation and LCAS that are used to group and maintain optical connections.

The PPE 15 includes a packet filter 65 receiving data packets as shown from the MAC 12. The packet filter 65 has an operative connection to packet filter tables 60 (actually a subset of all the packet filter tables as discussed above in relation to Fig. 4).

Packet filter 65 is the engine that processes the packets from data interfaces. The packet filter 65 is associated with and has access to packet filter table 60. For each incoming data packet, the packet filter 65 will extract data interface information and the packet's Layer-2 and/or MPLS headers, and use the packet filter table 60 to determine the encapsulation labels and the corresponding logical connection. The packet filter 65 forwards the packets into the corresponding optical connections so determined.

Packet filter 65 is connected to a packet forwarder 75 which is responsible for adding/stripping the labels, and forward packets to/from data and circuit interfaces.

Elements 65, 75, and 85 may be implemented any number of ways and with any number of physical elements such as logical entities within a software program. For high packet-switching performance, Elements 65, 75 and 85 can be implemented with specialize ASIC, FPGA, or off-the-shelf Network Processors. To satisfy pseudo-wire QoS requirements, further ASIC, FPGA and off-the-shelf Traffic Management chips may be required. Another example is a network processor unit complex which would include a network processing unit (NPU), memory, and optionally a traffic management chip with software coded to implement the invention running on the NPU. Another option would put all of these functions on one or more ASICs.

Packet forwarder 75 is also connected to a circuit filter 85 which has access to circuit filter table 80 (again, a subset of the circuit filter table maintained by the switch controller 20 as discussed above in relation to Fig. 4).

The circuit filter 85 is the engine that processes the packets coming from optical connections. Circuit filter 85 is associated with and has access to the circuit filter table 80. For each packet fetched from the optical connection, circuit filter 85

will extract the encapsulation label that identifies the data flow from the packet, and search the circuit filter table 80 for the outgoing data interface. If the packet is a control message (as determined by the identifiable encapsulation label for control messages), it will be forwarded to the switch controller 20 via the control message pathway as further shown in Fig. 7a. Otherwise, the circuit filter 85 strips off the label, and forwards the recovered packet to the corresponding data interface.

PPE controller 70 has a control connection to packet forwarder 75 and a control message pathway to switch controller 60. In addition, PPE controller 70 has access to session table 25 (again, a subset of the session filter table maintained by the switch controller 20 as discussed above in relation to Fig. 4).

The PPE Controller 70 is the logical entity that communicates with the switch controller 20. PPE controller 70 is associated with and has access to the session table 25, which maintains the mapping of control messages and outgoing optical connections. To inject a control message, PPE controller 70 searches the session table 25 to determine the encapsulating label and optical connection. Once the information is located, PPE controller 70 will encapsulate the control message and send out the control message via the optical connection (by way of the mapping engine 17, TDM switch fabric 30, and TDM line module 40).

The packet filter 65 and circuit filter 85 may be constructed as logical elements within a software program. As such these filters 65, 85 may share processing resources with the PPE controller 70 or may be separately constructed.

In more detail and as shown in Figure 7b, the packet filter table 60 has the following attributes:

- A Searching Key which includes the packet's (incoming) data interface and label information.
 - (Incoming) data interface: This is the interface that receives the packet. It can be the identification for either a physical or logical interface. The invention makes no assumption on how such information is actually obtained. However, the interface information is required for each packet being received.
 - Label: This can be, for example, a Layer-2 or MPLS header. A Layer-2 header can be an Ethernet MAC and VLAN tag, a Frame Relay DLCI, or an ATM VCI/VPI. It is noteworthy that a received packet may have

been encapsulated with a Layer-2 header and a MPLS label. In this case, two matching keys are defined: one with Layer-2 header; the other, MPLS label. In Figure 7b, Packet-Filter-1 and Packet-Filter-4 can be applied to the same packet.

- **Outgoing Optical Connection:** This is the connection that the packet will be injected into as it enters the provider network.
- **Encapsulation Label:** The label for each data flow. It will be encapsulated with the packet.
- **Filter Priority:** The importance of the filter. As mentioned above, a packet may be encapsulated with both Layer-2 and MPLS. Thus, two matching filters may be found. We use the Filter Priority to decide which filter should be applied to the packet. In Figure-7b, if a packet received from Port-1 that matches to both Packet-Filter-1 and Packet-Filter-4, Packet-Filter-1 will be chosen since it has a higher priority.
- **Guaranteed QoS:** This is an optional field when QoS (quality of service) is an issue. If so, each data flow should comply within a fixed traffic boundary. Otherwise, traffic congestion may result within an optical connection. This field maintains the guaranteed QoS for the flow. For packets that do not comply, a user-defined traffic conditioning mechanism will be used. The mechanism itself is beyond the scope of this invention.

As shown in Figure 7b, the packet filter table 60 is populated with data showing the various types of packets that may be processed including Ethernet, ATM, FR and MPLS. Indeed, in this populated packet filter table 60, the PPE 15 is handling 4 different flows each with a unique encapsulation label. The corresponding outgoing optical connection fields are associated with each of these packet types.

As further shown in Fig. 7a, the data packets that arrive at the packet filter may be in the form of a layer 2 data (L2) with an associated packet or frame structure encapsulating the L2 data. Alternatively, an MPLS data with an associated packet or frame structure may also arrive at the packet filter 65. At element 75, the element 75 pushes a pre-negotiated encapsulation label onto the L2 packet or MPLS packet. When a control message is received from switch controller 50 via PPE controller 70, the element 75 also pushes a pre-negotiated encapsulation label onto the control message. With the encapsulation label added, the data flow is next sent to the circuit

filter 85 before being output as a logical circuit (SNC or sub-network connection) to the next stage which is the mapping engine 17 as shown in Figure 4.

Fig. 15a shows in more detail the processing performed by the PPE 15 on a data packet received from a data interface. As shown therein, the PPE 15 receives (400) a packet from a data port and then the packet filter 65 parses (405) the layer-2 (and perhaps the MPLS header if present) and searches the packet filter table 60.

The packet filter 65 then decides (410) whether there is a match with the packet filter table 60. If not, then the packet is dropped (440) thereby ending processing for the received packet.

If there is a match, the flow proceeds and decides (415) if there is more than one matching filter which may be the case if the packet is encapsulated with both Layer-2 and MPLS headers (or other multiple headers as may be the case). More than one header causes the packet filter 65 to choose (445) the header with the highest priority (see filter priority field in Fib. 7b).

The traffic condition may then be determined (420). When a filter is found for a packet, the traffic condition for that flow, such as the bandwidth consumed by the flow, will be known. The packet filter 65 and packet filter table 60 keep track of the QoS information for all flows. If, by receiving this packet, it will cause the flow's QoS parameters (such as bandwidth consumption) to be over its defined limit, the PPE 15 will apply traffic conditioning to the packet, either dropping or tagging the packet. With this information, the packet filter 65 may then determine (425) if the traffic condition is within a QoS limit. The invention does not define the actual mechanism for the packet filter 65 to come to that decision 425; rather, it only operates on the final outcome. If not within the QoS limit, then the traffic condition or rule is followed 450 meaning that the traffic is dropped or tagged. If (455) not tagged, the packet is dropped (440). If it is tagged, the flow proceeds to the encapsulation (430) step. Steps 420, 425, 450, 455 are considered option and implemented only when QoS is a factor.

The encapsulation (430) involves looking up the encapsulation label from the packet filter table 60 and pushing the encapsulation label onto the packet as illustrated in Fig. 7a. Then, the encapsulated packet may be sent (435) out to the outgoing optical connection as defined in the packet filter table 60.

In general, the PPE 15 performs the following processes. Since each SONET cross-connection can carry traffic from multiple L2 users, it is necessary to be able to

distinguish individual user's frames at place where de-multiplexing takes place. The PPE takes care of this by pushing an encapsulation label onto every L2 frame that will enter the provider network. The encapsulation label may come from the negotiation between provider edges using LDP.

At exiting edge, the encapsulation label will be popped, and the original frames will be recovered and delivered out to the destination customer. This process is described below in more detail in relation to Fig. 15b and the circuit filter table of Fig. 7c.

Figure-7c: Circuit Filter Table

The Circuit Filter Table has the following attributes:

- Searching Key: Optical Connection and Label
 - Optical Connection: The connection where a packet is received. It can be a SONET VCG (Virtual Concatenation Group) or an optical interface
 - Label: This is the label that has been inserted at the ingress of the data flow. It is used to identify a specific data flow.
- Outgoing Data Interface: The interface where the packet to be forwarded. As shown in Figure-7c, all control messages go to "Host Interface", which is the Switch Controller in this case.
- Overwritten Label: It is possible that the customer may want to change a packet's Layer-2 label as it traverses through the optical network. One such instance is that the customers want to change Ethernet VLAN values to satisfy Ethernet bridging protocol requirements. Overwritten Label contains the new label information. PPE is responsible for the label over-writing.
- Guaranteed QoS: Each data flow must comply within a fixed traffic boundary. Otherwise, this may result in traffic congestion at outgoing data port. This field maintains the guaranteed QoS for the flow. For packets that do not comply, a user-defined traffic conditioning mechanism will be used. The mechanism itself is beyond the scope of this invention.

As shown in Fig. 15b, the PPE 15 performs the following processes when receiving a packet from an optical connection. First, the PPE fetches (500) the packet from an optical connection. The circuit filter 85 may then parse (505) the encapsulation label from the packet and use it to search the circuit filter table 80 (see Fig. 7c). The results of the circuit filter table 80 search are used to determine (510) if

there is exactly one match. If not, the packet is dropped (540) and this event is recorded.

If there is only one match, then the circuit filter 85 may determine (515) the traffic condition. Once again, the circuit filter is keeping track of the QoS parameters, (bandwidth, delay, and packet dropped etc.) for every flow. If by receiving this packet causes the flow's QoS parameters going over the limit, we will have to either drop or tag the packet.) The results of this determination (515) are used to decide (520) if the traffic condition is over the QoS limit. If yes, then the packet is (tagged or dropped) (545) according to the QoS rule stored in the circuit filter table 80 for that packet. A decision (550) is based on whether the packet is tagged or dropped: if to be dropped the flow proceeds to drop (540) the packet; otherwise, the flow proceeds to remove (525) the encapsulation label. Like the QoS processing described above in relation to Fig. 15a, these steps are option if QoS is not a factor in the system.

After removing (525) the label, the circuit filter 85 decides (530) whether to require overwriting of the packet header. See the description for the parameter above for details. If yes, the circuit filter 85 overwrites the header according to the entry for that circuit contained in the circuit filter table 80. If the entry indicates that the label is not to be overwritten than the PPE 15 sends out the packet through the data interface defined in the circuit filter table 80 for that packet. In this way, the data flow arriving from the provider network may be correctly routed to the correct data interface and, ultimately, to the correct client edge node.

Since the control messages come as labeled packets, the circuit filter table 80 will match them to "host interface". The sending step 535 will send regular packets to data interfaces, and control messages to this "host interface" which is the switch controller 20 itself.

Figure 16 and session table 7d further explain the control messaging procedures. PPE controller 70 implements the process of Fig. 16 with access to the session table 25 of Fig. 7d.

Figure-7d: Session Table

The Session Table has the following attributes:

- Searching Key: Control Message ID
 - Each control message carries a unique ID to identify which "peering session" it belongs to. A peering session is a logical connection

between two edge nodes. It is used to exchange control information between two nodes. For example, in pseudo-wire operation, the customer may apply LDP [RFC3036] to negotiate data flows. LDP operates over TCP. Between two edge nodes, all control messages go over a TCP session that can be uniquely identified with TCP Sender Port Number, and IP addresses. In this invention disclosure, we shall not specify the exact message ID format. However, it is reasonable to assume that each control message carries enough information to identify the session to which it belongs.

- As an example, in Figure-7d, there are three sessions that are identified with TCP and UDP port numbers.
- Outgoing Optical Connection: This is the connection that the control messages will be injected into.
- Encapsulation Label: The identifiable label for the control message. PPE will insert this label to the control message.
- Guaranteed QoS: All control messages within a session will have a fixed network resource level. This is designed to protect the control messages from potential congestion caused by regular data traffic.

The process begins by the PPE controller 70 receiving (600) a control message from the switch controller 20 which is then parsed (605) to find the ID as explained above.

The PPE controller 70 then searches (610) the session table 25 according to the control message ID parsed (605) from the control message. The results of the search are used to decide (615) if there is a match such that the corresponding entry may be retrieved from the session table 25. If not match, the message is dropped (640) and the event recorded. If there is a match, the PPE controller 70 may perform some QoS processing (steps 620, 625, 645, 650, 640) that are analogous to the QoS processing described above in relation to Figs. 15a and 15b such that a repetition here is not necessary. Again, this QoS processing is considered an optional but desirable feature.

After QoS processing, the PPE may then send (635) out the control message to the associated optical interface (identified by the entry in the session table 25 for that control message) as a data payload. Specifically, the control message is tunneled as payload within a SONET, SDH or OTN frame payload and thereby shares its fate with the packet data being carried by the provider network.

Provisioning of Pseudo-wires

The conventional LDP (label distribution protocol, RFC3036) is used by the invention to setup and manage pseudo wires: each pseudo-wire runs over a bi-directional cross-connection such as a SONET, SDH, or OTN cross-connection. Each pseudo-wire includes two unidirectional paths, one in each direction. Each provider edge initiates the setup of the path on behalf of ingress L2 traffic.

Each path may be uniquely identified by the triple <sender, receiver, encapsulation label>. The triple is part of the message sent between nodes during the label negotiation phase shown in Fig 17. The VCID is an example of an encapsulation label that may be used by the invention. A conventional VCID label is a 32-bit quantity that must be unique in the context of a single LDP session between two provider edges. For a given pseudo-wire, the same encapsulation label (e.g. VCID) must be used when setting up both paths.

As described during our discussion on Figure 3, to aggregate a data flow and thus establish a pseudo-wire, the network operator first downloads all the mapping information to the provider edge nodes. Through LDP, two provider edge nodes negotiate encapsulation label for a data flow.

To create a pseudo wire between two provider edges, the network operator needs to provide the IP addresses of the provider edges, and assign a, for example, 32-bit VCID to represent this pseudo wire. To support Ethernet VLAN services, the operator needs to feed VLAN-ID's to both provider edges as well.

Through LDP, two provider edge nodes exchange encapsulation label, physical port and VLAN information, and negotiate the encapsulation labels. Specifically, LDP will use Virtual Circuit FEC and Generic Label FEC during label negotiation. Upon completion, the provider edge nodes will program hardware for frame classification and MPLS label encapsulation. The detailed operation of LDP is conventional and beyond the scope of this invention.

Fig. 17 further explains the process of setting up a pseudo wire over optical network according to the invention. Essentially, Fig. 17 is a sequence diagram that performs the following processes.

1. Initially, there exists an operational optical connection between provider edge nodes (Node-1 and Node-2 in Fig. 17). Traditionally, in carrier networks, such connections are static in nature – they are not frequently modified once established.

2. Node-1 and Node-2 will establish a peering session over the optical connection. The method for session establishment is to inject control messages into the connection, and each control message is encapsulated with an identifiable label. (See the description for Figures 7 and 7d above)
3. Upon the establishment of the peering session, Network Operator will issue data flow setup requests to both Node-1 and Node-2. The request will include the following information:
 - a. The data interfaces that packets will traverse.
 - b. The optical connection the packets need to aggregate into.
 - c. The QoS (bandwidth) requirements for each flow.
 - d. Optionally, the packet Layer-2 label to be overwritten (see the description for Figure 7c)
4. The integrity of the requests is maintained by the network operators, and is beyond the scope of this invention.
5. Node-1 and Node-2 will exchange control messages and negotiate the labels to be used by the data flows. An example of the label negotiation is described in [draft-martini].
6. Upon the completion of the label negotiation, Node-1 and Node-2 will update the data-plane with the label information, that is, to populate the packet filter table 60 and the circuit filter table 80 on the PPE 15.

Data flow can now be transmitted over the optical connection.

Fig. 18 further explains the process of tearing down or deleting a pseudo wire according to the invention. Essentially, Fig. 18 is a sequence diagram that performs the following processes.

1. The Network Operator sends the deletion requests to both Node-1 and Node-2.
2. Node-1 and Node-2 will exchange control messages and withdraw the labels that are previously allocated for the data flow. In case of SONET connection failure or operational teardown, LDP is responsible for withdrawing labels at provider edges
3. Upon the completion of the operation, Node-1 and Node-2 will update the data plane by deleting the corresponding entries from the Packet/Circuit Filter Tables.

Fig. 19 further explains the process of handling outages on the optical connections that affect one or more pseudo wires according to the invention. Essentially, Fig. 19 is a sequence diagram that performs the following processes.

1. The optical connection between Node-1 and Node-2 is no longer working. This could be the result of a planned outage by the carriers, or a link failure in the network. The outage may be detected in any number of conventional fashions and such detection is outside the scope of this invention.
2. Node-1 and Node-2 will update the data-plane immediately. One action is to suspend all the relevant Packet/Circuit Filters on PPE. Another option is to reroute the traffic to another optical connection. The mechanism of rerouting at pseudo-wire level is beyond the scope of this invention.
3. Node-1 and Node-2 will notify the condition to Network Operator.

Alternative Architectures Benefiting From Invention

The switch fabric 32 is a generalized interconnect between line modules. The interconnects are for optical connections and may also include an additional packet flow interconnect to exchange packet data between modules prior to the mapping engine function. The implementation of the fabric interconnects is outside the scope the invention and does not impact the invention functions. Conceptually, it is convenient to consider two independent switch fabrics as shown in Figures 8 through 13b; the TDM switch fabric 30 for optical connections and the packet fabric 34 for packet data that has not been mapped to an optical connection. However, in practice the interconnect function may be implemented in any fashion and with any number of technologies. Examples of other fabric implementations include a single TDM switch fabric, a single packet switch fabric, and technologies may include any pure electrical, or a hybrid optical/electrical switch fabric.

Some higher-level architectural details and alternatives will be explored in this section. All of these architectures clearly benefit by utilizing the inventive concepts as further explained below.

The invention described herein may be implemented on any form of optical connection switch. Given the variety of sizes and designs of switches and the varying needs in data packet capacity requirements, it is natural that there are many possible

configurations for incorporating the functionality described in the invention into such switch designs.

Generally speaking, the functional elements of the switch described herein are not required to be oriented or arranged as shown in Fig 8. For example, the PPE 15 may be located on a dedicated field replaceable card independent of the line modules 40, switch controller 20, or switch fabric 32 as shown in Fig 9.

As further shown in the packet-data-enabled optical connection switch 5' configuration of Fig. 9, the packet server 90 contains the PPE 15 and mapping engine 17 while the MAC 12 is contained on a simplified Packet Access Line Module (PALM' 10'). The TDM Line Module 40 is a conventional optical connection originating/terminating module as in Fig 8. The switch 5' shown in Fig. 9 is a simplified diagram of a practical switch and has only one PALM' 10', one TDM line module 40 and one packet server 90 but it is to be understood that in a practical implementation that a plurality of these elements are included to provide the switch with greater capacity.

Comparing the Fig. 9 configuration of the switch 5' against the Fig. 8 configuration, the mapping engines 17 function identically. The PPE 15 functions are also identical but implementation would be different, thus PPE is labeled 15' in Fig. 9. The switch controller 20 and the tables 25, 60, 80 would also be the same other than differences in switch control coordination of flows to PPE and PPE to optical connection which is more complicated.

More specifically, the PPE 15' in Fig. 9 sends and receives traffic via the mapping engine 17 and packet fabric 34 while the PPE 15 in Fig. 8 may also send and receive traffic via a physical client port via the MAC 12. In both configurations the PPE's primary function is to manage pseudo-wires in optical connections and translate and manage packet data flows from/to the pseudo-wires.

In order to benefit from statistical multiplexing gain, many pseudo-wires (on the order of 1,000s or 10,000s) will be carried in each optical connection. The data flows that are translated into these pseudo-wires will normally connect to the packet-data-enabled optical connection switch over many different physical ports. These physical ports may be located on several different PALMs 10. The PPE 15 will aggregate these multiple pseudo-wires and use traffic shaping principals to share one or more optical connections between the pseudo-wires. The source/destination flow associated with each pseudo-wire may reach the PPE via a MAC 12 located on the PALM 10 with the PPE

15, or it may be forwarded via the packet fabric 34 from a PPE 15 located on another PALM 10. This is the architecture shown in Fig. 8.

As the space and power limits of the PALM 10 will limit the size and capacity of the PPE 15 that can be located on the PALM 10, it may be desirable to locate the PALM on a dedicated module like the packet server 90 shown in Fig. 9. In this configuration, the PPE' 15' operates as described above.

The packet server 90 is essentially another example of switch architecture with the PPE and other data functions included.

As described earlier, the implementation of the interconnect switch fabric 32 is beyond the scope of the invention. Depending on the implementation of the packet data interconnect function 34, it may be necessary to translate the packet data traffic from/to the PPE 15, 15' into a compatible format for the interconnect. In Fig 9, the packet fabric interface 16 is fulfilling this function. This is a detail that could be considered part of the packet fabric/interconnect implementation and removed from the figure as in Figs. 10-13a.

More specifically, the switch 5' may contain multiple packet server modules 55 to increase the packet processing capacity of the switch 5' and/or for redundancy as shown in Fig 10. As shown therein, n PALM' modules labeled 10'-1 through 10'-n are provided. In addition, j packet servers labeled 90-1 through 90-j are also provided.

Packet traffic transmitted between PALM' 10' cards and packet server 90 cards can be carried over a packet switch fabric 34 or interconnect as shown in Figure 10. The packet switch fabric 34 or interconnect may be implemented any number of ways. Examples of implementations include but are not limited to a dedicated packet switch element contained on a field replaceable switching card; dedicated backplane traces between PALMs 10' and packet servers 90; an asynchronous crossbar switch; or dedicated connections between PALMs 10' and packet servers 90 in the TDM switch fabric 30.

A packet switch fabric 34 or interconnect may be used in the packet-data-enabled optical connection switch 5' even if the architecture does not include packet server modules 90. As shown in Figure 8, a packet switch fabric 34 or interconnect can be used to transmit packets between PPEs located on multiple PALMs (e.g. between PPE 15-1 on PALM 10-1 and PPE 15-n on PALM 10-n. Transmitting packets between PPEs 15 in such a fashion allows aggregation of packet data from multiple physical ports that reside on different PALMs

An advantage of a packet-data-enabled optical connection switch 5, 5' is that the same network element can be used to switch a variety of types of traffic. Traditional TDM circuit traffic is switched similarly as on traditional optical connection switches via a TDM fabric such as TDM fabric 32 and TDM line modules 40, 41 as shown in Figure 12. Simultaneously, the packet-data-enabled optical connection switch 5, 5' can be switching L2 packet flows into pseudo-wires over optical connection as described in the invention and shown in Figure 13 for the case of a packet server architecture. The PPE 15' and PALM 10' may be implemented to also allow packet switching between packet data ports as shown in Figure 11.

As mentioned earlier, an intermediate provider node may have the capability to overwrite an encapsulation label. Such a node would most likely contain a PPE 15 or 15' and mapping engine 17 to perform this function. One reason to overwrite the encapsulation label at an intermediate node would be to aggregate multiple pseudo-wires arriving at the node on different optical connections onto a common outbound optical connection.

An example of the data path through packet-data-enabled optical connection switch with packet server architecture is shown in Figure 13a. In this example, an optical connection containing packet data traffic arrives at the switch on TDM line module 40-1 and is switched via the TDM switch fabric 32 to the mapping engine 17-1 located on packet server 90-1. The PPE 15-1 will process the recovered packet as described earlier but the outgoing data interface entry in the circuit filter table will contain a value that reserved for the PPE to loop the packet back into the PPE 15-1 similar to if it were to have arrived from the packet fabric/interconnect 34. The PPE 15-1 will then process the packet again and based on the packet filter table 60 send the packet to the mapping engine 17-1 to go out another optical connection. This other optical connection, originating from the mapping engine 17-1 is switched via the TDM switch fabric 32 to the associated outbound TDM LM, 40-m in figure 13a.

As noted previously, the different types of L2 traffic supported by the packet-data-enabled optical connection switch may require multiple MACs 12 and/or multiple types of PALMs 10, 10'. Additionally, the PALM 10, 10' may contain multiple physical ports that may or may not be sending/receiving the same type of L2 traffic.

In a general case, a sub-set of ports on the PALM may send/receive conventional TDM optical connection traffic so that the PALM also functions as a

TDM LM on a sub-set or all of the traffic. Similarly, a mixture of conventional TDM traffic and L2 traffic may arrive on the same physical port of a PALM. In this case, the L2 traffic is contained in a TDM transport frame that is multiplexed with other transport frames into a single high-speed TDM frame. In order to access the L2 traffic, the PALM 10, 10' would perform conventional TDM add/drop multiplexing (ADM) functionality to terminate the TDM connection containing the L2 traffic and pass the remaining TDM connections to the TDM switch fabric.

For example, a physical port on a PALM may be receiving/transmitting a SONET OC48 signal with the first 12 STSs carrying ATM traffic and the remaining 36 STSs carrying TDM circuit traffic that is to be switched to other TDM outbound ports on the switch. The PALM 10, 10' would first demultiplex the OC48 signal using conventional means. The resultant tributary that contained the ATM traffic would be terminated and the L2 packets recovered and forwarded to the PPE. The remaining TDM tributaries would be forwarded to the TDM switch fabric 32, similar to how they would have been handled had they arrived at the switch on a TDM LM port.

Example Of Inventive Operation

In this section, we walk through an example of how a carrier provisions a pseudo-wire between SONET switches, such as a CoreDirector® (CD) switch made by CIENA Corporation.

As shown in Figure 20, CD-1 (IP loopback address 1.1.1.1) and CD-2 (IP loopback 2.2.2.2) are provided in a network having other (unlabelled) CDs that serve as intermediate nodes in the provider network. A customer attaches to port 1 on CD-1 using VLAN ID 100, and to port 2 on CD-2 using VALN ID 200. Inside the SONET transport network, SNC-12 is established ahead of time. SNC-12 can be used to carry Ethernet traffic between CD-1 and CD-2.

Both CD-1 and CD-2 use LDP to discover each other. This allows both nodes to exchange control information to setup the pseudo wires. All control messages are tunneled through SNC-12 as SONET payload and encapsulated with a MPLS "IP4 Explicit NULL Label".

Once a SNC is in place, establishing a pseudo wire includes three basic steps:

- 1. Network Operator Provisioning:**

Each VCID uniquely identifies a pseudo wire between a pair of edge nodes. At each node we associate a port/VLAN with a remote edge (loopback address) and VCID. In the example, the network operator picks VCID 50 to identify the pseudo wire between (Port 1, VLAN 100) on CD-1 to (Port 2, VLAN 200) on CD-2. All necessary information is downloaded to CD-1 and CD-2.

2. MPLS Label Advertisement and Solicitation:

Upon the completion of the provisioning process, LDP automatically exchanges pseudo wire information between CD-1 and CD-2. CD-1 advertises MPLS label 1000 for VCID 50 to CD-2. Similarly, CD-2 advertises label 2000 for VCID 50 to CD-1.

3. Data Plane Setup:

After MPLS labels have been exchanged, the edge nodes program the data plane for pseudo-wire operation. CD-1 will program the PPE as follows:

- For all Ethernet frames received from Port 1 with VLAN 100, push label 2000, and send the frames through SNC-12.
- For all Ethernet frames carried over SONET arriving on SNC-12 with label 1000, rewrite VLAN-ID to 100, send them through Port 1.

Similar rules are configured on CD-2 for frames going to CD-1.

Advantages Of Invention:

- Martini's pseudo-wire approach provides a uniformed method to carry all types of layer-2 traffic over a carrier's backbone network. However, the backbone must be MPLS/IP-enabled. Traditionally, carriers are very careful

with setting up SONET cross-connections inside their networks. In many cases, SONET connections are well provisioned with a rich set of features for network resource allocation, traffic restoration, and link protection, etc. Thus, instead of building pseudo-wires over a MPLS backbone, it would be desirable to use SONET cross-connects to carry pseudo-wire traffic directly.

- If backbone networks deliver only layer-2 frames between edges, it may be more economical from both an equipment and management expense point of view to provide the "tunneling" functionality on top of the SONET cross-connects directly, rather than building another layer of tunneling mechanism running on top of optical transport networks.
- In the invention, optical transport networks can be used to support both **traditional voice traffic** as well as **data packets**. The transport backbones can be provisioned and administrated as they have been for years. Only at network edges, pseudo wires are established to transfer data traffic. Thus, the overall transport management system is not disturbed.
- By creating pseudo wires on top of SONET cross-connects, carriers can better utilize network resource by mapping individual user traffic onto SONET virtual concatenated trunks, and adapt mechanisms such as LCAS to fine-tune bandwidth reservations. Since the pseudo wires and the optical cross-connections are originated from the same edge nodes, this can potentially reduce network operation cost for carriers.
- The carriers can aggregate data traffic into transport networks directly from network edge. There is no need to introduce UNI or NNI interfaces to bring data traffic into the optical domain. Mapping pseudo-wires into pre-established SNC's automatically can eliminate the undesired effect of creating and deleting SNC's dynamically at user and network interfaces.
- From a hardware support point of view, this approach will leverage the scalable SONET switching capability in some of the SONET switches.

Carriers can bundle and aggregate pseudo-wires into fine-granular STS trunks.

It is important to realize that SONET STS trunks themselves are perfect for user flow isolation and bandwidth guarantees. Providing class-of-service or QoS at an STS granularity is hence a unique feature that routers cannot cheaply replace in the foreseeable future.

- This invention can aggregate both traditional Layer-2 as well as MPLS labeled traffic over optical transport networks. As a result, this invention can further help network providers to integration services, such as L2 and L2 VPN, more economically.

It is to be understood that the inventive concepts are not limited to SONET and also include SDH which is the prevailing standard in Europe and emerging standards such as OTN. In other words, although the invention is described mainly in terms of SONET in the interest of simplifying the description, the inventive concepts may be equivalently applied to SDH or OTN networks.

The invention being thus described, it will be obvious that the same may be varied in many ways. Such variations are not to be regarded as departure from the spirit and scope of the invention, and all such modifications as would be obvious to one skilled in the art are intended to be included within the scope of the following claims.